

 **PORTAL**
USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login
 Search: The ACM Digital Library The Guide SEARCH

THE ACM DIGITAL LIBRARY Feedback

+cached query* +reuse count*

Terms used: cached query reuse count

Sort results by

relevance

 Save results to a Binder

Display results

expanded form

 Open results in a new windowRefine these results with A
Try this search in The ACM

Results 1 - 20 of 127

Result page: 1 2 3 4 5 6 7 next >>

F

1 Simulation analysis of virtual geographic routing

David M. Nicol, Michael E. Goldsby, Michael M. Johnson

December 2004 WSC '04: Proceedings of the 36th conference on Winter simulation

Publisher: Winter Simulation Conference

Full text available:  pdf(461.36 KB) Additional Information: full citation, abstract, references, cited by

Bibliometrics: Downloads (6 Weeks): 4, Downloads (12 Months): 28, Citation Count: 1

Homeland defense applications will use large-scale ad-hoc networks of small devices. Routing is a crucial problem, for naive means do not scale well. Geographic Routing (GR) (Karp 2000; Giordano, Stojmenovic, and Blazevic 2003) offers hope for scalability, ...

2 A predicate-based caching scheme for client-server database architectures

Arthur M. Keller, Julie Basu

January 1996 The VLDB Journal — The International Journal on Very Large Data Bases,

Volume 5 Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(162.50 KB) Additional Information: full citation, abstract, references, cited by, index terms

Bibliometrics: Downloads (6 Weeks): 7, Downloads (12 Months): 51, Citation Count: 43

We propose a new client-side data-caching scheme for relational databases with a central server and multiple clients. Data are loaded into each client cache based on queries executed on the central database at the server. These queries are used to form ...

Keywords: Cache completeness, Cache currency, Caching, Multiple clients, Relational databases

3 Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren

November 1997 CASCON '97: Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research

Publisher: IBM Press

Full text available:  pdf(4.21 MB) Additional Information: full citation, abstract, references, index terms

Bibliometrics: Downloads (6 Weeks): 56, Downloads (12 Months): 761, Citation Count: 0

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event ...

4 Transparent caching with strong consistency in dynamic content web sites

 Cristiana Amza, Gokul Soundararajan, Emmanuel Cecchet
June 2005 ICS '05: Proceedings of the 19th annual international conference on Supercomputing
Publisher: ACM

Full text available:  pdf(370.07 KB) Additional Information: full citation, abstract, references

Bibliometrics: Downloads (6 Weeks): 6, Downloads (12 Months): 44, Citation Count: 0

We consider a cluster architecture in which dynamic content is generated by a database back-end and a collection of Web and application server front-ends. We study the effect of transparent query caching on the performance of such a cluster. Transparency ...

5 I/O reference behavior of production database workloads and the TPC benchmarks—an analysis at the logical level

 Windsor W. Hsu, Alan Jay Smith, Honesty C. Young
March 2001 ACM Transactions on Database Systems (TODS), Volume 26 Issue 1
Publisher: ACM

Full text available:  pdf(5.42 MB) Additional Information: full citation, abstract, references, cited by, index terms

Bibliometrics: Downloads (6 Weeks): 21, Downloads (12 Months): 144, Citation Count: 8

As improvements in processor performance continue to far outpace improvements in storage performance, I/O is increasingly the bottleneck in computer systems, especially in large database systems that manage huge amounts of data. The key to achieving ...

Keywords: I/O, TPC benchmarks, caching, locality, prefetching, production database workloads, reference behavior, sequentiality, workload characterization

6 Managing multiple and distributed ontologies on the Semantic Web

A. Maedche, B. Motik, L. Stojanovic
November 2003 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 12 Issue 4

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(375.18 KB) Additional Information: full citation, abstract, references, cited by, index terms

Bibliometrics: Downloads (6 Weeks): 29, Downloads (12 Months): 216, Citation Count: 11

In traditional software systems, significant attention is devoted to keeping modules well separated and coherent with respect to functionality, thus ensuring that changes in the system are localized to a handful of modules. Reuse is seen as the key method ...

Keywords: Multiple and distributed ontologies, Ontology evolution

7 Efficient use of the query optimizer for automated physical design

Stratos Papadomanolakis, Debabrata Dash, Anastasia Ailamaki
September 2007 VLDB '07: Proceedings of the 33rd international conference on Very large data

bases

Publisher: VLDB Endowment

Full text available:  pdf(2.24 MB) Additional Information: full citation, abstract, references

Bibliometrics: Downloads (6 Weeks): 14, Downloads (12 Months): 55, Citation Count: 0

State-of-the-art database design tools rely on the query optimizer for comparing between physical design alternatives. Although it provides an appropriate cost model for physical design, query optimization is a computationally expensive process. The ...

8 Just fast keying: Key agreement in a hostile internet William Aiello, Steven M. Bellovin, Matt Blaze, Ran Canetti, John Ioannidis, Angelos D.

Keromytis, Omer Reingold

May 2004 ACM Transactions on Information and System Security (TISSEC), Volume 7 Issue 2

Publisher: ACM

Full text available:  pdf(324.39 KB) Additional Information: full citation, abstract, references, cited by, index terms

Bibliometrics: Downloads (6 Weeks): 23, Downloads (12 Months): 156, Citation Count: 4

We describe Just Fast Keying (JFK), a new key-exchange protocol, primarily designed for use in the IP security architecture. It is simple, efficient, and secure; we sketch a proof of the latter property. JFK also has a number of novel engineering parameters ...

Keywords: Cryptography, denial-of-service attacks**9 Context-based prefetch – an optimization for implementing objects on relations**

Philip A. Bernstein, Shankar Pal, David Shutt

December 2000 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 9 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(142.24 KB) Additional Information: full citation, abstract, references, index terms

Bibliometrics: Downloads (6 Weeks): 4, Downloads (12 Months): 65, Citation Count: 0

When implementing persistent objects on a relational database, a major performance issue is prefetching data to minimize the number of round-trips to the database. This is especially hard with navigational applications, since future accesses are unpredictable. ...

Keywords: Caching, Object-oriented database, Object-relational mapping, Prefetch**10 DEW: DNS-enhanced web for faster content delivery** Balachander Krishnamurthy, Richard Liston, Michael Rabinovich

May 2003 WWW '03: Proceedings of the 12th international conference on World Wide Web

Publisher: ACM

Full text available:  pdf(331.13 KB) Additional Information: full citation, abstract, references, cited by, index terms

Bibliometrics: Downloads (6 Weeks): 3, Downloads (12 Months): 48, Citation Count: 2

With a key component of latency on the Web being connection set up between clients and Web servers, several ways to avoid connections have been explored. While the work in recent years on Content Distribution Networks (CDNs) have moved some content 'closer' ...

11 ARCube: supporting ranking aggregate queries in partially materialized data cubes

Tianyi Wu, Dong Xin, Jiawei Han
June 2008 SIGMOD '08: Proceedings of the 2008 ACM SIGMOD international conference on Management of data
Publisher: ACM

Full text available:  pdf(411.77 KB) Additional Information: full citation, abstract, references, index terms

Bibliometrics: Downloads (6 Weeks): 30, Downloads (12 Months): 30, Citation Count: 0

Supporting ranking queries in database systems has been a popular research topic recently. However, there is a lack of study on supporting ranking queries in data warehouses where ranking is on multidimensional aggregates instead of on measures of base ...

Keywords: data cube, partial materialization, ranking aggregate queries

12 Labels and event processes in the Asbestos operating system

Steve Vandevooght, Petros Efstathopoulos, Eddie Kohler, Maxwell Krohn, Cliff Frey, David Ziegler, Frans Kaashoek, Robert Morris, David Mazières
December 2007 ACM Transactions on Computer Systems (TOCS), Volume 25 Issue 4
Publisher: ACM

Full text available:  pdf(1.21 MB) Additional Information: full citation, abstract, references, index terms

Bibliometrics: Downloads (6 Weeks): 34, Downloads (12 Months): 210, Citation Count: 0

Asbestos, a new operating system, provides novel labeling and isolation mechanisms that help contain the effects of exploitable software flaws. Applications can express a wide range of policies with Asbestos's kernel-enforced labels, including controls ...

Keywords: Information flow, labels, mandatory access control, process abstractions, secure Web servers

13 Summary cache: a scalable wide-area web cache sharing protocol

Li Fan, Pei Cao, Jussara Almeida, Andrei Z. Broder
June 2000 IEEE/ACM Transactions on Networking (TON), Volume 8 Issue 3
Publisher: IEEE Press

Full text available:  pdf(220.29 KB) Additional Information: full citation, references, cited by, index terms

Bibliometrics: Downloads (6 Weeks): 25, Downloads (12 Months): 446, Citation Count: 76

Keywords: ICP, Web cache, Web proxy, bloom filter, cache sharing

14 Exploring the tradeoff between performance and data freshness in database-driven Web servers

Alexandros Labrinidis, Nick Roussopoulos
September 2004 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 13 Issue 3
Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(387.89 KB) Additional Information: full citation, abstract, references, index terms

Bibliometrics: Downloads (6 Weeks): 8, Downloads (12 Months): 65, Citation Count: 0

Personalization, advertising, and the sheer volume of online data generate a staggering amount of dynamic Web content. In addition to Web caching, view materialization has been shown to accelerate the generation of dynamic Web content. View materialization ...

15 Cooperative scans: dynamic bandwidth sharing in a DBMS
Marcin Zukowski, Sándor Héman, Niels Nes, Peter Boncz
September 2007 VLDB '07: Proceedings of the 33rd international conference on Very large data bases
Publisher: VLDB Endowment
Full text available: pdf(482,51 KB) Additional Information: full citation, abstract, references
Bibliometrics: Downloads (6 Weeks): 8, Downloads (12 Months): 50, Citation Count: 0
This paper analyzes the performance of concurrent (index) scan operations in both record (NSM/PAX) and column (DSM) disk storage models and shows that existing scheduling policies do not fully exploit data-sharing opportunities and therefore result in ...

16 Queue: Volume 6 Issue 3
 May 2008 issue Volume 6 Issue 3
Publisher: ACM
Full text available: pdf(5,56 MB) digital edition Additional Information: full citation
Bibliometrics: Downloads (6 Weeks): 0, Downloads (12 Months): 0, Citation Count: 0

17 QPipe: a simultaneously pipelined relational query engine
 Stavros Harizopoulos, Vladislav Shkapenyuk, Anastassia Ailamaki
June 2005 SIGMOD '05: Proceedings of the 2005 ACM SIGMOD international conference on Management of data
Publisher: ACM
Full text available: pdf(506,36 KB) Additional Information: full citation, abstract, references, cited by
Bibliometrics: Downloads (6 Weeks): 15, Downloads (12 Months): 132, Citation Count: 3
Relational DBMS typically execute concurrent queries independently by invoking a set of operator instances for each query. To exploit common data retrievals and computation in concurrent queries, researchers have proposed a wealth of techniques, ranging ...

18 Rendering with coherent layers
Jed Lengyel, John Snyder
August 1997 SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques
Publisher: ACM Press/Addison-Wesley Publishing Co.
Full text available: pdf(1,32 MB) Additional Information: full citation, references, cited by, index terms
Bibliometrics: Downloads (6 Weeks): 4, Downloads (12 Months): 51, Citation Count: 11
Keywords: Talisman, affine transformation, image compositing, image-based rendering, sprite

19 The holodeck ray cache: an interactive rendering system for global illumination in nondiffuse environments
Gregory Ward, Maryann Simmons

October 1999 ACM Transactions on Graphics (TOG) . Volume 18 Issue 4

Publisher: ACM

Full text available:  pdf(935.74 KB)

Additional Information: full citation, abstract, references, cited by, index terms

Bibliometrics: Downloads (6 Weeks): 5, Downloads (12 Months): 66, Citation Count: 9

We present a new method for rendering complex environments using interactive, progressive, view-independent, parallel ray tracing. A four-dimensional holodeck data structure serves as a rendering target and caching mechanism for interactive ...

Keywords: illumination, image reconstruction, mesh generation, ray tracing, rendering system, virtual reality

20 An approximate global illumination system for computer generated films

 Eric Tabellion, Arnaud Lamotte

August 2004 SIGGRAPH '04: ACM SIGGRAPH 2004 Papers

Publisher: ACM

Full text available:  pdf(819.51 KB)

Additional Information: full citation, abstract, references, cited by, index terms

Bibliometrics: Downloads (6 Weeks): 30, Downloads (12 Months): 179, Citation Count: 7

Lighting models used in the production of computer generated feature animation have to be flexible, easy to control, and efficient to compute. Global illumination techniques do not lend themselves easily to flexibility, ease of use, or speed, and have ...

Keywords: distributed ray tracing, global illumination, irradiance caching, micro-polygon, rendering

Results 1 - 20 of 127

Result page: 1 2 3 4 5 6 7 next >>

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2008 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  Adobe Acrobat  QuickTime  Windows Media Player  Real Player